

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In Re The Application of:)
John B. Duffie III et al.)
Serial No.: 10/087,376)
Filed: March 1, 2002)
For: METHOD TO OPTIMIZE THE)
LOAD BALANCING OF PAR-)
ALLEL CO-PROCESSORS)
)

Examiner: Daftuar, Saket K.

Art Unit: 2151

Cesari and McKenna, LLP
88 Black Falcon Avenue
Boston, MA 02210
June 25, 2007

By EFS

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

AMENDED APPEAL BRIEF

In response to the Notice of Appeal filed Dec. 19, 2006, and the Notice of Panel Decision from Pre-Appeal Brief Review mailed February 2, 2007 and the Notification of Non-Compliant Appeal Brief mailed May 24, 2007, the Applicants hereby submits this Amended Appeal Brief.

TABLE OF CONTENTS

I. REAL PARTY IN INTEREST.....	page 3
II. RELATED APPEALS AND INTERFERENCES.....	page 4
III. STATUS OF CLAIMS.....	page 5
IV. STATUS OF AMENDMENTS.....	page 6
V. SUMMARY OF CLAIMED SUBJECT MATTER.....	page 7
VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL.....	page 10
VII. ARGUMENT.....	page 11
VIII. CONCLUSION.....	page 21
VIV. CLAIM APPENDIX.....	page 22

I. REAL PARTY IN INTEREST

The real party in interest is Cisco Technology, Inc. by an Assignment recorded at reel 012672 frame 0284.

II. RELATED APPEALS AND INTERFERENCES

The Applicants and the Applicants' legal representatives know of no related appeals or interferences that will directly affect, be directly affected by, or have a bearing on the Board's decision in the present appeal.

III. STATUS OF CLAIMS

The status of the claims is:

A. Total Number of Claims in Application

Claims 1-35 stand pending in the Application.

B. Status of All the Claims

No claims stand cancelled.

Claims 1-35 stand rejected.

No claims are withdrawn from consideration on appeal.

No claims stand allowed.

C. Claims on Appeal

Claims 1-35 are currently being appealed.

IV. STATUS OF AMENDMENTS

An Amendment after Final under 37 C.F.R. §1.116 was filed on November 20, 2006, subsequent to the Final Rejection dated Sept. 19, 2006. The Amendment was entered by the Examiner, as indicated in the Advisory Action mailed Dec. 5, 2006.

V. SUMMARY OF CLAIMED SUBJECT MATTER

This summary describes exemplary embodiments of the invention.

Independent claim 1 is directed to a method for selecting a coprocessor from a plurality of coprocessors to process a packet. See Application page 9, lines 1-3. A size of a packet is determined. See page 4, lines 17-18 and page 10, lines 9-11. In response to the size of the packet, a cost associated with the with the packet is determined. See page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5. The cost represents a load associated with processing the packet. See page 10, lines 8-9. An anticipated load for each coprocessor in the plurality of coprocessors is determined using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. A coprocessor from the plurality of coprocessors is selected based upon the anticipated load. See page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 17 is directed to an apparatus for selecting a coprocessor from a plurality of coprocessors to process a packet. See page 9, lines 1-3. A memory (see page 7, lines 24-30 and Fig. 3, 340) contains one or more software routines configured to determine a size of a packet (see page 4, lines 17-18, page 10, lines 9-11) and to determine a cost associated with the packet in response to the size of the packet (see page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5). The cost represents a load associated with processing the packet. See page 10, lines 8-9. A processor (see page 7, lines 4-11 and Fig. 3, 320) is configured to execute the software routines to determine an anticipated load for each coprocessor in the plurality of coprocessors using the cost (see page 10, line 26 to page 11, lines 6 and Fig. 6, 604) and to select a coprocessor from the plurality of coprocessors based on the anticipated load (see page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618).

Independent claim 21 is directed to an intermediate device configured to select a coprocessor from a plurality of coprocessors to process a packet. See page 9, lines 1-3. The device includes means for determining a size of the packet (see page 4, lines 17-18, page 10, lines 9-11) and for determining a cost associated with the packet in response to the size of the packet (see page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5).

The cost represents a load associated with processing the packet. See page 10, lines 8-9. The device includes means for determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. Further, the device includes means for selecting a coprocessor based upon the anticipated load. See page 11, lines 7-11 and page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 22 is directed to a computer readable media comprising computer executable instruction for execution in a processor for selecting a coprocessor from a plurality of coprocessors to process a packet. See page 9, lines 1-3. The instruction are provided for determining a size of the packet (see page 4, lines 17-18, page 10, lines 9-11) and determining a cost associated with the packet in response to the size of the packet (see page 9, lines 15 to page 10, line 13 and Fig. 4, 402 and Fig. 5). The cost represents a load associated with processing the packet. See page 10, lines 8-9. Instructions are also provided for determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. Further, instructions are provided for selecting a coprocessor from the plurality of coprocessors based upon the anticipated load. See page 11, lines 7-11 and page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 23 is directed to a method for selecting a processor for processing a packet. See page 9, lines 1-3. A size of a packet is determined. See page 4, lines 17-18 and page 10, lines 9-11. A cost associated with the packet of that size is determined. See page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5. The cost represents a load associated with processing the packet. See page 10, lines 8-9. An anticipated load for the processor is determined using the cost. See page 10, line 26 to page 11, lines 6 and Fig. 6, 604. The processor is selected based upon the anticipated load. See page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

Independent claim 26 is directed to a method for selecting a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet. See page 9, lines 1-3. A cumulative load for each processor is determined, the cumulative

load representing load due to packets currently awaiting processing. See page 10, line 28 to page 11, line 4. A cost for processing the received packet at each coprocessor is determined, at least in part, in response to the size of the received packet and a processing rate of that coprocessor. See page 9, line 15 to page 10, line 13, Fig. 4, 402 and Fig. 5. The cumulative load is combined with the cost at each coprocessor to create an anticipated load for each coprocessor. See page 10, line 26 to page 11, line 6. The anticipated loads of all the coprocessors are compared. See page 11, lines 6-11. In response to the comparison, a particular coprocessor of the plurality of coprocessors is selected to perform the processing operation on the received packet. See page 12, line 15-21 and Fig. 6, 616, and 618.

Independent claim 31 is directed to an apparatus to select a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet. See page 9, lines 1-3. A plurality of queues are configured to store packets currently awaiting processing, each queue associated with one of the coprocessors. See page 7, lines 19-23 and Fig. 3, 323a, 323b, 323c. Each queue is associated with a cumulative load that represents a load to process packets in that queue. See page 10, line 28 to page 11, line 4. A data structure is configured to store processing rates, each processing rate associated with one of the coprocessors. See col. 9, lines 23-30 and Fig. 3, 345. A processor is configured to determine a size of the received packet (see page 4, lines 17-18 and page 10, lines 9-11), and in response to the size of the received packet, and the processing rate of each coprocessor, determine a cost to perform a processing operation on the received packet of each coprocessor (see page 9, line 15 to page 10, line 13 and Fig. 4, 402 and Fig. 5). The processor is further configured to combine the cost at each coprocessor with the cumulative load at that coprocessor to create an anticipated load at each coprocessor (see page 10, line 26 to page 11, line 6), and to select a particular coprocessor to perform the processing operation on the received packet in response to comparison of the anticipated load at each processor (see page 11, lines 7-11, page 12, lines 15-21 and Fig. 6, 606, 616, and 618.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

- A. Whether claims 1-35 are unpatentable under 35 U.S.C. §103 over U. S. Patent No. 6,587,866 to Modi et al. (hereinafter “Modi”), in view of U.S. Patent No. 5,613,069 to Walker (hereinafter “Walker”) where the references fail to teach or suggest several elements of the claims and there is no motivation or evidentiary basis to support combination of the references.

VII. ARGUMENT

A. Grouping of Claims

For each ground of rejection that the Applicants contests herein that applies to more than one claim, such additional claims, to the extent separately identified and argued below, do not stand or fall together.

B. Legal Standard

In rejecting claims under 35 U.S.C. §103, the Examiner bears the initial burden of presenting a prima facie case of obviousness. See, e.g., In re Rijckaert, 9 F.3d 1531, 1532 (Fed. Cir. 1993). To establish a prima facie case of obviousness, the references, when considered in their entirety, must teach or suggest all of the claimed limitations. If the references fail to teach or suggest any one of the claimed limitations, then the rejection should be reversed. An Examiner may not, moreover, resort to speculation, unfounded assumption or hindsight reconstruction to supply deficiencies in the references. In re Warner, 379 F.2d 1011, 1017 (CCPA 1967). In addition, the Examiner must set forth specific reasons why one skilled in the art would be motivated to select the cited features for combination in the manner claimed. In re Rouffet, 149 F.3d 1350, 1357 (Fed. Cir. 1998). Such a requirement is intended to avoid impermissible hindsight, in which the inventor's own disclosure is used as a blueprint for piecing together the prior art to defeat patentability. In re Dembiczak, 175 F.3d 994, 999 (Fed. Cir. 1999).

C. Rejection under 35 U.S.C. §103(a) over Modi in view of Walker

1. Claims 1, 6 and 8-25

Independent claim 1, representative of claims 1, 6 and 8-25, recites (with emphasis added):

1. A method for selecting a coprocessor from a plurality of coprocessors to process a packet, the method comprising steps of:
 - determining a size of the packet;
 - determining a cost associated with the packet in response to the size of the packet, the cost representing a load associated with processing the packet;*

*determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost; and
selecting the coprocessor from the plurality of coprocessors based on the anticipated load.*

Modi, U.S. Patent No. 6,587,866, discloses a load-balancing scheme for distributing packets among a plurality of server nodes in a clustered processing system. See Modi col. 3, lines 18-22 and col. 7, lines 61-64. The scheme uses one of a plurality of load balancing policy types, included “non-affinity” policy types where packets are distributed to any server node, and “affinity” policy types where packets from a single client are sent to the same server node. See col. 7, lines 61 to col. 8, lines 2 and col. 2, lines 42-47. The policy types include the well-known “round robin” and “weighted round robin” techniques. In a weighted round robin technique, nodes are assigned differing weights to account for differing processing capabilities, and accordingly the probability of a particular packet being sent to a particular node is dependent on that node’s relative weight. Modi does this by giving some nodes more entries in a packet distribution table (PDT) 304 than others. Specifically, “a high performance sever node is given more entities in PDT 304 than a slower server node that is able to able process less traffic. In this way, the high-performance server node will, on average, receive more traffic than the slower server node.” See col. 8, lines 21-25 and col. 9, lines 46-52. To implement “affinity” for nodes, the IP address of a received packet is hashed over the packet distribution table (PDT 304) and in response to the hash the packet placed in a “bucket” associated with a particular node. See col. 10, lines 32-35 and col. 14, lines 59-67.

Walker, U.S. Patent No. 5,613,069, discloses a packet switching network in which an “average transmission delay” (t_t) is computed using formulas that include “average packet size,” abbreviated b_{av} . For example, Walker states $t_t = (n * b_{av}) / s$ where n is a number of packets and s is a transmission rate. See Walker abstract and col. 18, lines 7-67. Walker also discusses that any “software delay” should be assumed “approximately constant.” See col. 18, lines 50-55.

i. The References, if Combined, Do Not Teach or Suggest All Limitations of the Claims.

The combination of Walker and Modi fails to teach or suggest all the claim limitations of claim 1. First, neither reference teaches or suggests "determining a cost associated with the packet in response to the *size of the packet*, the cost representing a load associated with processing the packet" and using such a cost in selecting a coprocessor from the plurality of coprocessors to process the packet.

As the Applicant describes in the background section of the Application, prior techniques have suffered shortcomings since they simply balance **the numbers** of packets received by coprocessors, and generally do not take into consideration the amount of resources that may be required to process any particular packet, which is may be related to **packet size**. For example, as described at page 3, lines 22-29 of the Specification, three 100-byte packets may be processed more rapidly than two 1400-byte packets, yet many prior techniques, such as the well known round-robin and shortest queue first (SQF) techniques, largely ignore this, instead merely counting numbers of packets. Thus, with many prior techniques, a coprocessor processing three 100-byte packets would be incorrectly considered more burdened than the one processing two 1400-byte packets.

The techniques discussed in Modi suffer the very shortcomings discussed by the Applicant in the background section of the Application. Modi does not consider packet size in determining which node (coprocessor) should process a packet. Instead, Modi simply balances **the number of packets to be processed**, for example by using a round-robin and/or hashing the **IP address of packets** over a packet distribution table. See Modi col. 9, lines 28-32, col. 10, lines 32-33 and col. 14, lines 59-67. A fast node in Modi may get more (a greater number of) packets to process, while a slow node may only get a few packets to process. However, the size of each packet is never considered.

In the Final Office Action of Sept. 19, 2006, the Examiner writes "examiner interprets determining which bucket to chose based on IP address packet size and when the user sends data, the data is divided into packets discloses determining the size of the

packet.” The Applicant respectfully traverses these interpretations as they distort the teachings of Modi by “interpreting” described quantities into clearly different quantities. “IP address” is a generally unique identifier of source or destination and differs considerably from “packet size” which is a measure of an amount of data, for example a number of bytes of data. Similarly, the act of dividing data into packets differs considerably from determining the size of a given packet. One skilled in the art would not interpret the Modi reference in the manner proposed by the Examiner, and accordingly would not reach the conclusions proposed by the Examiner.

In an attempt to address the manifest shortcomings of Modi, the Final Office Action combines it with Walker (a combination that the Applicant believes is improper as discussed below). However, combination with Walker does not remedy the shortcomings of Modi, as Walker also lacks any teaching or suggestion of the aspects of the Applicant’s claims that Modi lacks. Specifically, Walker does not suggest the Applicant’s claimed “determining a cost associated with the packet in response to the size of the packet.”

Walker simply discusses using an “average packet size” in calculating a transmission delay across a large packet switch network, for example by dividing average packet size by a transmission rate. See Walker abstract and col. 18, lines 7-67. The transmission delay Walker describes differs sharply from the claimed cost associated with processing a packet, as it in no way measures processing of a packet. Indeed, Walker instead suggests that any cost associated with processing a packet (in a sense software delay) should be assumed constant, and, as such, is not determined in response to the size of the packet. See Modi col. 18, lines 50-55.

Second, neither Modi nor Walker teaches or suggests “determining an anticipated load for each coprocessor in the plurality of coprocessors using the cost,” nor using the anticipated load in selecting a coprocessor. The Applicant novelly considers an anticipated load that a coprocessor would be subject to if it were to process the packet at issue, in deciding if the coprocessor should process the packet. In contrast, Modi simply dis-

cusses static load-balancing techniques, for example, a weighted round-robin and/or hash operation, which are not responsive to actual anticipated loads, nor take into account the cost of processing a particular packet. For example, a node may be assigned a weight in Modi, yet this weight is not determined in response to the cost of processing a particular packet at issue, it is merely some predetermined figure.

Similarly, there is no teaching or suggestion of these features in Walker, nor does Final Office Action of Sept 19, 2006 allege they are present there.

ii. There is No Motivation or Evidentiary Basis to Support Combination of the References

The Applicant further urges there is no motivation or evidentiary basis to combine Modi and Walker. As well expressed in MPEP 2143.01(III) “[t]he mere fact that references can be combined or modified does not render the resultant combination obvious unless the prior art also suggests the desirability of the combination.” See In re Rouffet, 149 F.3d 1350, 1357 (Fed. Cir. 1998). Some rationale must be expressly or impliedly contained in the prior art, or reasoned from knowledge generally available to one of ordinary skill in the art. See *Id.*

Neither the Final Office Action dated Sept. 19, 2006, nor the Advisory Action dated Dec. 5, 2006, cite to any expressed or implied suggestion or motivation in the references themselves. Further, neither the Final Office Action dated Sept. 19, 2006, nor the Advisory Action dated Dec. 5, 2006, present a reasoned argument why there is some suggestion or motivation in the knowledge of one of ordinary skill in the art. Instead, all that is presented is the conclusory statement:

Therefore, it would have bee [sic] obvious to one ordinary skilled in the art at the time the invention was made to securely controlled [sic] dynamically [sic] switching and distributing packets to the server based on their packet size and the cost associated with the packets.

See Office Action dated Sept. 19, 2006, page 3-4.

In contrast, there is persuasive evidence in the references that motivates against combination. If Modi were somehow modified to look to “packet size,” rather than to “IP address,” as suggest by the Examiner, Modi would be rendered unsatisfactory for one of its intended purposes. “[I]f proposed modification would render the prior art invention being modified unsatisfactory for its intended purpose, then there is no suggestion or motivation to make the proposed modification.” See MPEP §2143.01 and In re Gordon, 733 F.2d 900, (Fed. Cir. 1984). An important purpose of Modi is to be able to provide “client affinity.” Such purpose is expressed throughout the reference, including in the title, which reads “Method for Distributing Packets to Server Nodes Using Network *Client Affinity* and Packet Distribution Table” (emphasis added). If one were to somehow substitute the mention of “packet size” from Walker into Modi in place of “IP address,” as the Examiner suggests, Modi would likely no longer be able to implement client affinity, and thus would be unsuitable for one of its primary intended purposes.

Accordingly, for the above reasons, rejection of claims 1, 6, and 8-25 should be reversed.

Claims 2 and 3

Dependent claim 2, representative of claims 2 and 3 recites:

2. The method of claim 1 wherein the step of determining a cost further comprises the step of:
calculating the cost using a rate associated with processing the packet.

In the Final Office Action of Sept. 19, 2006, the Examiner cites only to Modi col. 1, at lines 38-39 in relation to this claim element, and makes no allegation that Walker has any relevance thereto. The cited portion of Modi reads:

The present invention relates to clustered computer systems with multiple nodes that provides services in a scalable manner.

Modi at col. 1, lines 38-39 has little to do with “calculating the cost [representing a load associated with processing the packet] using a rate associated with processing the

packet.” Rather, the cited passage of Modi simply expresses a desire that a computer system be expandable (scalable) so that other components may be added to increase capabilities.

Further, even if one were to look to Walker in connection with this aspect of the claims, they would find little teaching or suggestion thereof. Walker describes assuming cost associated with processing a packet (in a sense software delay) is a constant, and thus discourages one from actively calculating it using a rate. See Walker col. 18, lines 50-55.

Accordingly, for the above reasons, rejection of claims 2 and 3 should be reversed.

Claims 4 and 5

Dependent claim 4, representative of claims 4 and 5, recites:

4. The method of claim 2 wherein the step of calculating the cost further comprising the step of:
dividing the packet's size by the rate.

In the Final Office Action of Sept. 19, 2006, the Examiner cites only to Modi col. 15, lines 39-41 in relation to this claim element, and makes no allegation that Walker has any relevance thereto. The cited portion of Modi reads:

When the user sends data, a connection is established, the data is divided into packets, which are sent to the server during the connection.

Modi at lines 39-41 has little to do with dividing a packet's size by a rate associated with processing the packet as part of a step of calculating cost. Instead, Modi at col. 15, lines 39-41 deals with “dividing” data in the sense of segmenting data (i.e. breaking a large chunk of data into several smaller chunks of data).

Further, even if one were to look to Walker in connection with this aspect of the claims, they would find little teaching or suggestion thereof. Walker describes assuming cost associated with processing a packet (in a sense software delay) is a constant, and

thus discourages one from actively calculating it by dividing a packet size by a rate associated with processing the packet. See Walker col. 18, lines 50-55.

Accordingly, for the above reasons, rejection of claims 4 and 5 should be reversed.

Claim 7

Dependent claim 7 recites:

7. The method of claim 1 wherein the step of determining an anticipated load further comprises the step of:

adding the cost to a cumulative load associated with each coprocessor in the plurality of coprocessors.

In the Final Office Action of Sept. 19, 2006, the Examiner cites only to Modi at col. 1, lines 60-65 in relation to this claim element, and makes no allegation that Walker has any relevance thereto. The cited portion of Modi reads:

It is desirable for such a system to be efficient in order to accommodate as much traffic as possible with a minimal amount of response time. It is desirable for a system to be “scalable,” so that additional server nodes can be added and balancing between the nodes can be modifiable to provide a serve as demand for the system increase.

Modi at col. 1, lines 60-65 simply discusses the general goal that a system should be “efficient” and “scalable.” As such, it lacks any teaching or suggestion of a cost, representing a load associated with processing the packet, being added to a cumulative load on a coprocessor, as part of a step of determining an anticipated load for the coprocessor. Further, even if one were to look to Walker in connection with this aspect of the claims, they would find little teaching or suggestion thereof. Walker lacks any disclosure of cumulative loads associated with coprocessors of a plurality of coprocessors, or of adding a cost thereto.

Accordingly, for the above reasons, rejection of claim 7 should be reversed.

Claims 26-35

Independent claim 26, representative of claims 26-35, recites (with emphasis added):

26. A method for selecting a coprocessor from a plurality of coprocessors to perform a processing operation on a received packet, the method comprising steps of:

determining a cumulative load for each coprocessor, the cumulative load representing load due to packets currently awaiting processing at that coprocessor;

determining a size of the received packet;
determining a cost for processing the received packet at each coprocessor, the cost determined, at least in part, in response to the size of the received packet and a processing rate of that coprocessor;

combining the cumulative load and the cost at each coprocessor, to create an anticipated load for each coprocessor;

comparing the anticipated loads of all the coprocessors; and
selecting, in response to the comparing, a particular coprocessor of the plurality of coprocessors to perform the processing operation on the received packet.

The combination of Modi and Walker fails to teach or suggest all the claim limitations of claim 26. To the extent that the above discussions are relevant to various limitations of claim 26, in the interest of brevity, the reader is referred thereto. The Applicants would like to emphasize, however, that both Modi and Walker lack any teaching or suggestion of "determining a cumulative load for each coprocessor, the cumulative load representing load due to packets currently awaiting processing at that coprocessor" and "combining the cumulative load and the cost at each coprocessor, to create an anticipated load for each coprocessor."

As discussed above, Modi simply discloses static load balancing techniques, for example a weighted round-robin and/or hash operation, that are not responsive to actual cumulative loads due to packets currently awaiting processing, nor take into account the actual cost of processing the packet at issue. For example, a node may be assigned a weight in Modi, yet this weight is not determined in response to cumulative loads, nor the cost of processing the packet, it is merely some predetermined static figure. Similarly, Walker lacks any disclosure of relating to this topic.

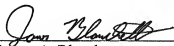
Accordingly, as the rejections of claims 26-35 should be reversed.

VIII. CONCLUSION

The Applicants respectfully submits that the claims are allowable over the art of record. Accordingly, the Applicants requests that the rejection of all claims be reversed.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,



James A. Blanchette
Reg. No. 51,477
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
(617) 951-2500

VIV. CLAIMS APPENDIX

- 1 1. (PREVIOUSLY PRESENTED) A method for selecting a coprocessor from a plural-
2 ity of coprocessors to process a packet, the method comprising steps of:
3 determining a size of the packet;
4 determining a cost associated with the packet in response to the size of the packet,
5 the cost representing a load associated with processing the packet;
6 determining an anticipated load for each coprocessor in the plurality of coproces-
7 sors using the cost; and
8 selecting the coprocessor from the plurality of coprocessors based on the antici-
9 pated load.
- 1 2. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of determin-
2 ing a cost further comprises the step of:
3 calculating the cost using a rate associated with processing the packet.
- 1 3. (ORIGINAL) The method of claim 2 wherein the rate is stored in a lookup table.
- 1 4. (ORIGINAL) The method of claim 2 wherein the step of calculating the cost further
2 comprising the step of:
3 dividing the packet's size by the rate.

1 5. (PREVIOUSLY PRESENTED) The method of claim 2 wherein the step of calculat-
2 ing the cost further comprises the step of:
3 multiplying the packet's size by a multiplicative inverse of the rate.

1 6. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of determin-
2 ing a cost further comprises the step of:
3 applying the packet's size to a lookup table containing one or more cost values to
4 determine the cost.

1 7. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of determin-
2 ing an anticipated load further comprises the step of:
3 adding the cost to a cumulative load associated with each coprocessor in the plu-
4 rality of coprocessors.

1 8. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of selecting
2 the coprocessor further comprises the step of:
3 selecting the coprocessor from a group of one or more coprocessors whose anti-
4 cipated load is a minimum load.

1 9. (ORIGINAL) The method of claim 8 wherein the coprocessor is selected using a
2 scheduling algorithm.

1 10. (PREVIOUSLY PRESENTED) The method of claim 1 wherein the step of selecting
2 the coprocessor further comprises the step of:
3 determining if a port associated with the packet is congested.

1 11. (PREVIOUSLY PRESENTED) The method of claim 10 wherein the step of select-
2 ing the coprocessor further comprises the step of:
3 selecting the coprocessor from a group of one or more coprocessors whose anti-
4 pated load is not a minimum load.

1 12. (PREVIOUSLY PRESENTED) The method of claim 10 wherein the step of select-
2 ing the coprocessor further comprises the step of:
3 selecting the coprocessor from a group of one or more coprocessors whose anti-
4 pated load is a minimum load.

1 13. (ORIGINAL) The method of claim 1 further comprising the step of:
2 incrementing a cumulative load associated with the selected coprocessor.

1 14. (PREVIOUSLY PRESENTED) The method of claim 13 wherein the step of incre-
2 menting a cumulative load further comprises the step of:
3 adding the cost to the cumulative load.

1 15. (ORIGINAL) The method of claim 1 further comprising the step of:

2 decrementing a cumulative load associated with the selected coprocessor.

1 16. (PREVIOUSLY PRESENTED) The method of claim 15 wherein the step of decre-
2 menting a cumulative load further comprises the step of:
3 subtracting the cost from the cumulative load.

1 17. (PREVIOUSLY PRESENTED) An apparatus for selecting a coprocessor from a
2 plurality of coprocessors to process a packet, the apparatus comprising:
3 a memory containing one or more software routines, including a software routine
4 configured to determine a size of the packet, and to determine a cost associated with the
5 packet in response to the size of the packet, the cost representing a load associated with
6 processing the packet; and
7 a processor configured to execute the software routines to determine an antici-
8 pated load for each coprocessor in the plurality of coprocessors using the cost and to se-
9 lect the coprocessor from the plurality of coprocessors based on the anticipated load.

1 18. (ORIGINAL) The apparatus of claim 17 further comprising:
2 a data structure;
3 wherein the cost is determined using information contained in the data structure.

1 19. (ORIGINAL) The apparatus of claim 18 wherein the information contained in the
2 data structure includes the cost.

1 20. (ORIGINAL) The apparatus of claim 18 wherein the information contained in the
2 data structure includes a rate the coprocessor can process the packet.

1 21. (PREVIOUSLY PRESENTED) An intermediate device configured to select a co-
2 processor from a plurality of coprocessors to process a packet, the intermediate device
3 comprising:

4 means for determining a size of the packet, and for determining a cost associated
5 with the packet in response to the size of the packet, the cost representing a load associ-
6 ated with processing the packet;

7 means for determining an anticipated load for each coprocessor in the plurality of
8 coprocessors using the cost; and

9 means for selecting the coprocessor based on the anticipated load.

1 22. (PREVIOUSLY PRESENTED) A computer readable media comprising computer
2 executable instructions for execution in a processor for selecting a coprocessor from a
3 plurality of coprocessors to process a packet, the instructions for:

4 determining a size of the packet, and determining a cost associated with the
5 packet in response to the size of the packet, the cost representing a load associated with
6 processing the packet;

7 determining an anticipated load for each coprocessor in the plurality of coproces-
8 sors using the cost; and

9 selecting the coprocessor from the plurality of coprocessors based on the anti-
10 pated load.

1 23. (PREVIOUSLY PRESENTED) A method for selecting a processor for processing a
2 packet, the method comprising steps of:
3 determining a size of the packet;
4 determining a cost associated with the packet of that size, the cost representing a
5 load associated with processing the packet;
6 determining an anticipated load for the processor using the cost of the packet if
7 processed by the processor;; and
8 selecting the processor based on the anticipated load.

1 24. (PREVIOUSLY PRESENTED) The method of claim 23 wherein the step of deter-
2 mining a cost comprises the step of:
3 calculating the cost using a rate associated with processing of the packet; and
4 wherein the rate is stored in a lookup table.

1 25. (PREVIOUSLY PRESENTED) The method of claim 23 wherein the step of deter-
2 mining a cost further comprises the step of:
3 applying the size of the packet to a lookup table containing cost values associated
4 with particular sizes.

1 26. (PREVIOUSLY PRESENTED) A method for selecting a coprocessor from a plural-
2 ity of coprocessors to perform a processing operation on a received packet, the method
3 comprising steps of:
4 determining a cumulative load for each coprocessor, the cumulative load repre-
5 senting load due to packets currently awaiting processing at that coprocessor;
6 determining a size of the received packet;
7 determining a cost for processing the received packet at each coprocessor, the cost
8 determined, at least in part, in response to the size of the received packet and a processing
9 rate of that coprocessor;
10 combining the cumulative load and the cost at each coprocessor, to create an an-
11 ticipated load for each coprocessor;
12 comparing the anticipated loads of all the coprocessors; and
13 selecting, in response to the comparing, a particular coprocessor of the plurality of
14 coprocessors to perform the processing operation on the received packet.

1 27. (PREVIOUSLY PRESENTED) The method of claim 26, wherein the step of select-
2 ing further comprises the step of:
3 selecting a coprocessor with minimum anticipated load to perform the processing
4 operation on the received packet.

1 28. (PREVIOUSLY PRESENTED) The method of claim 26, further comprising the step
2 of:

3 determining if congestion is present at an output port associated with the received
4 packet, and if congestion is present, selecting a coprocessor with non-minimum anti-
5 pated load to perform the processing operation on the received packet.

1 29. (PREVIOUSLY PRESENTED) The method of claim 26, wherein the step of deter-
2 mining a cumulative load for each coprocessor further comprises the step of:

3 determining, for each coprocessor, sizes of the packets currently awaiting proc-
4 essing at that coprocessor and using the sizes in conjunction with the processing rate of
5 that coprocessor to determine the cumulative load.

1 30. (PREVIOUSLY PRESENTED) The method of claim 26 wherein the processing op-
2 eration is an encryption operation.

1 31. (PREVIOUSLY PRESENTED) An apparatus to select a coprocessor from a plural-
2 ity of coprocessors to perform a processing operation on a received packet, the apparatus
3 comprising:

4 a plurality of queues configured to store packets currently awaiting processing,
5 each queue associated with one of the coprocessors, each queue associated with a cumu-
6 lative load that represents a load to process packets in that queue;

7 a data structure configured to store processing rates, each processing rate associ-
8 ated with one of the coprocessors; and
9 a processor configured to determine a size of the received packet, and in response
10 to the size of the received packet, and the processing rate of each coprocessor, determine
11 a cost to perform a processing operation on the received packet at each coprocessor, the
12 processor further configured to combine the cost at each coprocessor with the cumulative
13 load at that coprocessor to create an anticipated load at each coprocessor, and to select a
14 particular coprocessor to perform the processing operation on the received packet in re-
15 sponse to comparison of the anticipated load at each coprocessor.

1 32. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the processor is
2 further configured to select a coprocessor with minimum anticipated load to perform the
3 processing operation on the received packet.

1 33. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the processor is
2 further configured to determine if congestion is present at an output port associated with
3 the received packet, and if congestion is present, select a coprocessor with non-minimum
4 anticipated load to perform the processing operation on the received packet.

1 34. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the cumulative
2 load associated with each coprocessor is determined in response to sizes of packets
3 awaiting processing in the queue associated with that coprocessor and the processing rate
4 of that coprocessor.

1 35. (PREVIOUSLY PRESENTED) The apparatus of claim 31, wherein the processing
2 operation is an encryption operation.

EVIDENCE APPENDIX

None.

RELATED PROCEEDINGS APPENDIX

None.